

Smart Courts and Digital Justice: A Governance Framework for Responsible Artificial Intelligence in Judicial Decision-Making

Anurag Sharma¹, Ms. Kritika Kakkar², Dr. Neha Sharma³, Dr. Akansha Pandey⁴, Dr Parul⁵

¹Assistant Professor, School of Legal Studies, K. R. Mangalam University, Gurugram, Haryana.

Email id: sanurag986@gmail.com

Orcid Id: <https://orcid.org/0000-0003-1581-3357>

²Assistant Professor, School of Legal Studies, K. R. Mangalam University, Gurugram, Haryana

Email id: kritika.03jan@gmail.com

³Assistant Professor, School of Legal Studies, K. R. Mangalam University, Gurugram, Haryana

Email id: nsharma.sbp@gmail.com

Orcid Id - <https://orcid.org/0000-0003-2503-4074>

⁴Assistant Professor, School of Legal Studies, K. R. Mangalam University, Gurugram, Haryana

Email id: dr.akanshapandey01@gmail.com

Orcid id - <https://orcid.org/0009-0002-7140-9922>

⁵Assistant Professor, School of Legal Studies, K. R. Mangalam University, Gurugram, Haryana

Email: parul69sethi11@gmail.com

HOW TO CITE:

Anurag Sharma, Kritika Kakkar, Neha Sharma, Akansha Pandey, Parul (2026). Smart Courts and Digital Justice: A Governance Framework for Responsible Artificial Intelligence in Judicial Decision-Making. International Journal of Special Education, 41(9s), 379-391.

COPYRIGHT STATEMENT:

Copyright: © 2026 Authors.
Open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).

ABSTRACT:

Artificial intelligence (AI) is now used in many parts of digital court systems, including e-filing, registry scrutiny, transcription, translation, legal research, document review, and case-management support. This review examines how AI can help smart courts while protecting the values that make adjudication legitimate: legality, fairness, reason-giving, accountability, judicial independence, and public trust. The study uses a qualitative scoping-review design. It draws on peer-reviewed legal scholarship, empirical studies, technical literature, institutional frameworks, regulatory materials, and official court-technology documents published mainly between 2018 and 2026. The review shows that AI is most appropriately deployed in administrative, informational, and assistive tasks that improve speed, access, language inclusion, and document management. Uses that predict outcomes, influence bail or sentencing, recommend risk scores, or draft judicial reasons create higher risks, including bias, opacity, privacy intrusion, automation bias, and diffusion of responsibility. The article proposes a human-centred and rights-based governance framework in which AI operates as accountable court infrastructure rather than an autonomous adjudicator. Smart courts therefore require not only technology, but also clear policies, independent validation, audit trails, contestability, data-protection safeguards, vendor controls, and judicial capacity building. AI can strengthen digital justice only when it supports, and does not replace, human judicial reasoning.

Keywords: smart courts; artificial intelligence; digital justice; judicial decision-making; algorithmic accountability; legal technology; access to justice; AI governance

Introduction

Courts across the world are becoming more digital. Many now use e-filing, digital case records, virtual hearings, automated notifications, searchable judgment

databases, translation tools, and data dashboards. Within this wider shift, AI has become a practical tool rather than a distant possibility. It can classify files, retrieve precedents, transcribe hearings, translate judgments, summarise long records, and support workflow management. The European Commission for the Efficiency of Justice has recognised that AI may improve the efficiency and quality of justice, but only if it is used consistently with fundamental rights, non-discrimination, quality, security, transparency, and user control [1].

The key question is not whether courts should use technology. Courts already depend on digital systems. The more difficult question is how far AI should enter the judicial sphere. Judicial decision-making is different from ordinary administrative processing because it affects liberty, property, dignity, family status, criminal responsibility, and constitutional rights. Soudin explains that AI may support, reshape, or even substitute parts of judicial work, but each option raises different legitimacy concerns [2]. Re and Solow-Niederman similarly argue that AI adjudication can affect important court values, including consistency, individualised justice, transparency, and institutional authority [3].

Empirical studies also show that AI should not be treated as automatically superior to human judgment. Machine-learning systems may improve some prediction tasks, including pretrial risk prediction, under controlled conditions [4]. At the same time, studies of recidivism tools have questioned whether proprietary risk-assessment systems perform better than simple human or statistical baselines [5]. Other research shows that human users may interpret algorithmic advice differently across groups, which can reproduce unfairness even when the tool is formally advisory [6]. These findings matter because an AI output can influence a judge even if it is not formally binding. Generative AI has made these

questions more urgent. GenAI tools can summarise files, draft text, translate documents, and retrieve authorities. They can also produce fabricated citations, inaccurate summaries, and confident but legally unsound reasoning. Recent scholarship on GenAI at the bench therefore stresses both its promise and its risks [7]. International frameworks, including the OECD AI Principles, UNESCO Recommendation on the Ethics of AI, and the EU AI Act, also point toward human oversight, transparency, accountability, safety, and rights protection as minimum conditions for responsible AI [8-10]. India provides a useful example of both promise and caution. Official sources describe AI-supported tools such as SUPACE for factual-matrix understanding and precedent search, as well as transcription, translation, e-

filing scrutiny, SUVAS-based judgment translation, and the National Judicial Data Grid [11-13]. These tools can improve access and efficiency. However, digital infrastructure does not by itself guarantee fair adjudication. Courts need governance rules that define what AI may do, what it must not do, and who remains accountable.

This article asks: How can AI be integrated into smart courts in a way that improves efficiency and access to justice while preserving fairness, transparency, accountability, and human judicial responsibility? Its contribution is to bring together legal theory, empirical evidence on human-algorithm interaction, AI governance principles, and court-administration practice into one court-specific framework. The article argues that judicial AI should be designed as accountable decision support, not as a substitute for human adjudication.

Literature Review

Literature on AI and courts has developed in three broad phases. The first phase focused on digital justice and online courts as ways to reduce delay, cost, and procedural complexity. The second phase examined predictive analytics and computational decision support, especially in bail, sentencing, parole, and outcome prediction. The third and most recent phase examines GenAI,

large language models, legal agent systems, and the governance of AI assistance in adjudication.

A central theme is the tension between efficiency and legitimacy. CEPEJ remains important because it treats judicial AI as a rights issue, not only a performance issue. Its principles of fundamental rights, non-discrimination, quality and security, transparency, fairness, and user control provide a baseline for responsible judicial AI [1]. The OECD and UNESCO frameworks also connect AI governance to human dignity, democratic values, privacy, transparency, fairness, and human oversight [8,9]. The NIST AI Risk Management Framework strengthens this approach by focusing on governance, mapping, measurement, and management of AI risks across the lifecycle of a system [14].

Legal scholarship has questioned whether AI can perform the justificatory role of courts. Sourdin distinguishes between AI that supports judicial work and AI that replaces judicial tasks, warning that dignity, responsiveness, and contextual discretion cannot be reduced to computation [2]. Re and Solow-Niederman frame AI adjudication as a challenge to the values that courts are expected to serve [3]. Floridi and Cowls also offer a useful ethical framework based on beneficence, non-maleficence, autonomy, justice, and explicability [15].

Empirical evidence is more cautious than many policy claims. Kleinberg et al. show that machine prediction may improve certain pretrial decisions under specific modelling assumptions [4]. However, Dressel and Farid show that the COMPAS recidivism tool had important limits in accuracy and fairness [5]. Green and Chen further show that adding an algorithm to a decision process does not automatically make the result better or fairer [6]. Courts should therefore avoid equating prediction with adjudication. Prediction estimates a future event; adjudication requires public reasons based on law, facts, procedure, and context.

Recent studies on bias, legal agent systems, and procedural fairness add further caution. Krištofik argues that AI-supported judicial decision-making can reproduce older institutional biases through

new technology [16]. Technical work such as AgentsCourt shows that legal agent systems can simulate court debate and use legal knowledge augmentation, but it also brings legal information processing closer to decision-like outputs [17]. Završnik warns that big-data criminal justice may weaken case-specific narratives [18]. Lopes argues that AI in adjudication can affect perceptions of procedural fairness even when it is introduced as a response to human bias [19].

The GenAI literature adds a new risk profile. Socol de la Osa and Remolina show that GenAI can support access and efficiency but can also misinform judicial reasoning through hallucinations, opacity, and weak accountability [7]. Sourdin and Barry also note that courts perform many functions beyond deciding disputes, including administrative, communicative, and public-trust functions [20].

The research gap is clear. Existing work often examines technical performance, ethical principles, or court digitisation separately. Fewer studies integrate these strands into a practical governance model that separates low-risk assistive AI from high-risk adjudicatory AI and links each category to specific controls. This article addresses that gap by proposing a court-specific governance framework grounded in legal theory, empirical evidence, institutional practice, and AI-risk governance.

Research Objectives

1. To identify major AI applications relevant to smart courts and digital justice systems.
2. To distinguish administrative, informational, analytical, and adjudicatory uses of AI in courts.
3. To assess the benefits of AI for speed, efficiency, language inclusion, legal information access, and court management.
4. To analyse the risks of AI-supported legal decision-making, especially bias, opacity, privacy intrusion, hallucination, and automation bias.
5. To evaluate legal and ethical safeguards required for accountable, transparent, and rights-compatible AI use in courts.

6. To propose a governance framework in which AI supports judicial work without replacing human adjudication.

Materials and Methods / Methodology

This study uses a qualitative scoping-review design. It is doctrinal, analytical, and socio-technical in orientation. The review examines legal scholarship, empirical studies, technical AI literature, institutional reports, court-technology documents, and AI governance frameworks. It does not conduct a survey, experiment, or statistical meta-analysis. This design is appropriate because the research question is concerned with conceptual integration, legal evaluation, and governance design rather than with estimating a pooled numerical effect.

The review followed a structured but non-quantitative source-selection process. Searches were planned around major academic and legal research sources, including Scopus, Web of Science, HeinOnline, Google Scholar, SSRN, publisher databases, official court websites, government portals, and international governance bodies. Priority was given to material published between 2018 and 2026 because judicial AI has changed quickly during this period. Foundational sources outside this period were retained when they remained important to the legal or governance debate.

Search terms included combinations of artificial intelligence, smart courts, digital justice, judicial decision-making, AI in courts, algorithmic bias, risk assessment, generative AI, legal decision support, court administration, transparency,

accountability, access to justice, human oversight, and AI governance. Sources were screened first for relevance to courts or judicial decision-making, then for credibility, and finally for their contribution to the themes of applications, benefits, risks, safeguards, or governance design. The final synthesis relied on 20 core sources cited in the reference list, supported by institutional and regulatory materials.

Inclusion criteria were: direct relevance to courts, judicial decision-making, or court administration; publication by a peer-reviewed journal, academic publisher, recognised standards body, court, government, or international organisation; and substantive contribution to benefits, risks, governance, or legal theory. Exclusion criteria were: vendor marketing, anonymous blogs, duplicated commentary, sources without identifiable authorship or institutional responsibility, and materials focused on AI uses outside justice systems.

The analytical procedure used thematic synthesis. First, sources were read for claims about AI applications in courts. Second, these claims were coded under five themes: applications, benefits, risks, ethical-legal issues, and governance safeguards. Third, the themes were compared with AI governance frameworks and empirical evidence on human-algorithm interaction. Fourth, the findings were organised into tables and figures. Because this is a qualitative scoping review, analytical adequacy was assessed through triangulation across legal theory, empirical evidence, technical literature, and institutional materials rather than statistical testing.

Table 1. Structured search and screening strategy for the qualitative scoping review

Review component	Operational approach	Purpose
Academic databases	Scopus, Web of Science, HeinOnline, Google Scholar, SSRN	To identify peer-reviewed legal, empirical, and interdisciplinary scholarship
Institutional sources	CEPEJ, OECD, UNESCO, NIST, official court and government sources	To identify governance standards and official court-technology practice

Screening focus	Relevance to AI in courts, judicial decision-making, legal accountability, or court administration	To exclude general AI materials with weak justice-system relevance
Inclusion criteria	Reliable authorship, institutional responsibility, legal or empirical relevance, and substantive governance value	To improve credibility and reduce unsupported claims
Exclusion criteria	Vendor marketing, anonymous commentary, duplicated news items, and unrelated technology sources	To avoid promotional or weakly supported material
Analytical synthesis	Thematic coding of applications, benefits, risks, safeguards, and governance controls	To connect technical capability with court-specific legal values

Results

The review produced four main findings. First, AI applications in courts are diverse and do not carry the same legal risk. Second, the strongest case for AI concerns administrative and informational support, not autonomous adjudication. Third, the main risks are

socio-technical. They arise from the interaction between data, models, institutional incentives, legal culture, vendors, and human users. Fourth, responsible smart courts require governance that is legal, technical, procedural, and organisational at the same time.

Table 2. AI applications across the judicial process and required safeguards

Judicial stage	AI application	Main value	Risk level and required control
Pre-filing and public access	Chatbots, plain-language guidance, eligibility triage	Improves user orientation and access to legal information	Low-medium; disclaimers, human helpdesk, accessibility review
Filing and registry scrutiny	OCR, metadata extraction, e-filing defect detection	Reduces procedural friction and filing delays	Low-medium; registry officer validation and audit trail
Case management	Listing support, clustering similar cases, workload dashboards	Improves scheduling and resource allocation	Medium; prevent opaque pressure on judges or parties
Legal research	Precedent search, retrieval-augmented summarisation	Accelerates identification of relevant authorities	Medium-high; citation verification and disclosure rules
Transcription and translation	ASR, speech-to-text, multilingual judgment translation	Improves record accuracy and language access	Medium; human review for legal meaning and evidentiary nuance

Document and evidence review	Entity extraction, document grouping, anomaly detection	Supports management of large case records	Medium-high; data minimisation, privilege and confidentiality controls
Risk or outcome prediction	Bail, sentencing, recidivism, case-outcome prediction	May support structured assessment	High; independent validation, contestability, and strict limits
Judgment drafting support	Summaries, draft outlines, legal issue maps	Can reduce drafting burden	High; judge-owned reasoning, verified citations, and no autonomous reasons

Table 2. AI applications across the judicial process and required safeguards

The classification shows that AI is most appropriate where it improves process management without deciding rights. Registry defect detection, metadata extraction, translation,

transcription, and legal information retrieval are usually defensible when they remain auditable and subject to human validation. By contrast, risk scoring, outcome prediction, and autonomous drafting of reasons are high-risk because they may influence legal reasoning, distort party equality, or introduce opaque assumptions.




Figure 1. AI-supported judicial process flowchart showing assistive AI entry points and human judicial control.

Figure 1 shows where AI can safely support the judicial workflow: filing support, registry scrutiny, scheduling, legal research, transcription, translation, and document organisation. The decisive tasks--fact evaluation, legal interpretation,

discretion, reason-giving, and final judgment--remain under human judicial authority. This separation is important because administrative efficiency must not become adjudicatory delegation.

Table 3. Benefits of AI in courts and governance conditions for legitimacy

Benefit	Court function improved	Countervailing risk	Governance condition
Speed	Faster scrutiny, retrieval, transcription, and translation	Fast but inaccurate processing may harm rights	Human verification and quality thresholds
Efficiency	Reduced repetitive administrative work	Automation may hide errors or shift work to litigants	Process mapping and measurable service standards
Access to justice	Plain-language support and multilingual access	Chatbots may give misleading legal information	Disclaimers, referral pathways, and periodic review
Consistency	Uniform metadata, defect identification, and classification	Rigid categories may ignore context	Appeal or correction channel for affected users
Transparency	Dashboards and searchable records	Surveillance or performance pressure on judges	Purpose limitation and independence safeguards
Knowledge management	Improved retrieval from judgments and case files	Hallucinated or outdated authorities	Source-linked retrieval and mandatory citation checking

The benefits of AI are real but conditional. AI can reduce delay, support language access, and improve retrieval across large legal records. The same tools can undermine justice if their outputs are inaccurate, biased, hidden, or treated as automatically correct. In courts, efficiency is legitimate only when it is tied to accountability and review.

Proposed Governance Model for Responsible Judicial AI

The proposed model treats AI as part of governed court infrastructure, not as an isolated software tool. User interfaces, court applications, AI services, data repositories, and governance controls must be designed together. This means that identity management, encryption, access control, validation, audit logs, disclosure rules, and human oversight should be built into the system from the beginning.

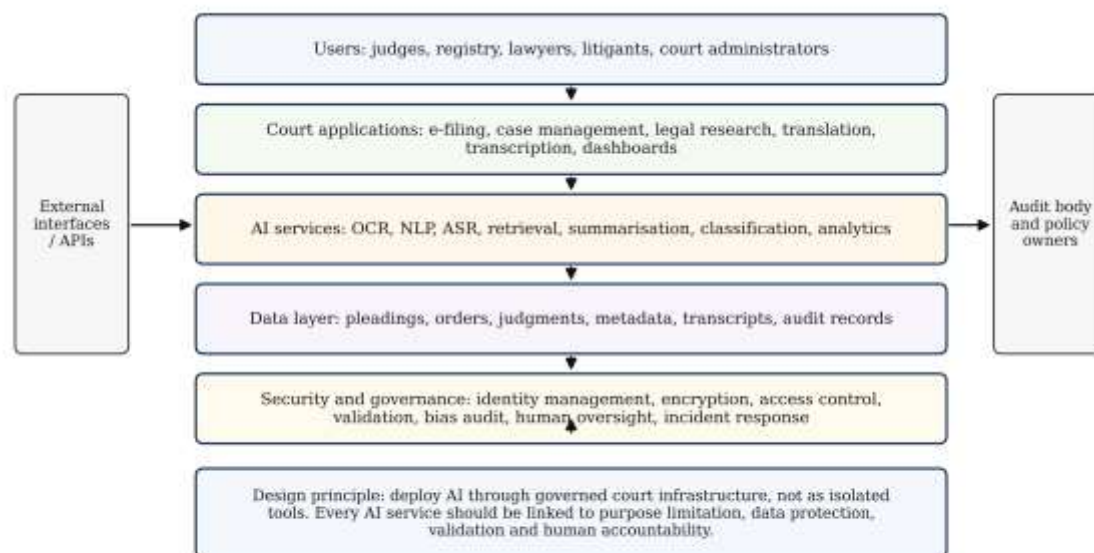


Figure 2. Architecture for responsible AI integration in court systems.

Figure 2 presents a layered architecture. AI services sit between court applications and the data layer, but they are constrained by security and governance controls. External interfaces and

policy owners should not be afterthoughts. They are part of the accountability structure that allows courts to monitor how AI tools are used.

Table 4. Ethical and legal issues with operational safeguards

Issue	Risk in court context	Operational safeguard	Primary accountability owner
Fairness	Discriminatory outputs or unequal impact on vulnerable groups	Bias audits, equality-impact assessment, representative evaluation	Court administration and oversight body
Transparency	Parties cannot know or challenge AI influence	Disclosure of AI use, model cards, output logs	Court and authorised system owner
Privacy	Exposure of sensitive case data, witness data, or privileged material	Data minimisation, encryption, access control, retention limits	Court data-protection officer
Accountability	Responsibility diffused among judges, vendors, and technologists	Named responsible officers, procurement duties, audit records	Judicial institution

Contestability	Litigants cannot challenge algorithmic assumptions	Correction, explanation, and objection mechanisms	Presiding court and registry
Due process	AI output influences reasons without party participation	No hidden adjudicatory use; party notice where rights are affected	Judge and court administration
Reliability	Hallucinations, inaccurate summaries, or obsolete law	Benchmark testing, source verification, incident reporting	System owner and judicial user
Vendor governance	Dependence on proprietary tools and unclear data processing	Contractual audit rights, data localisation where required, exit plans	Procurement authority

The safeguard matrix shows that high-level ethical principles are not enough. Courts need operational controls: procurement standards,

documentation, bias testing, impact assessment, incident reporting, audit trails, training, and review mechanisms.

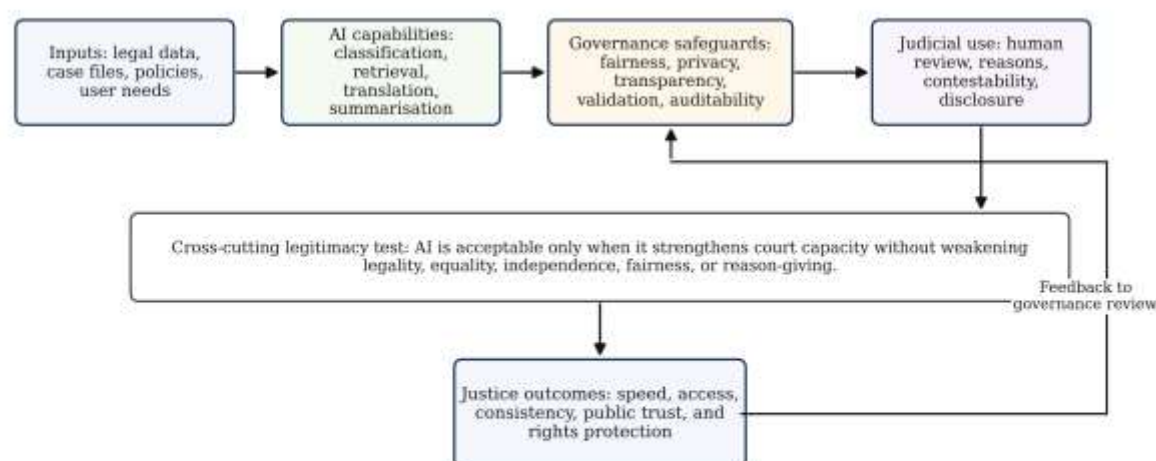


Figure 3. Conceptual framework for responsible AI-assisted digital justice.

Figure 3 summarises the article's central argument. AI systems use legal data and technical capabilities, but they become legitimate only when filtered through safeguards and human judicial use. Justice outcomes then feed back into governance review. A smart court is therefore not simply a court with AI tools. It is a court where

digital systems are governed so that efficiency remains subordinate to legality, fairness, and reasoned judgment.

Discussion

The findings support a careful view of AI in courts. AI is useful when it solves bounded

administrative or informational problems. It becomes risky when users treat mathematical-looking outputs as neutral legal judgment. This distinction is consistent with CEPEJ's user-control principle and UNESCO's emphasis on human oversight and human dignity [1,9]. In courts, human oversight must mean more than a judge approving an output after it has already shaped the case. It must include the ability to understand, question, reject, and explain the role of AI in the reasoning process.

Algorithmic accuracy is not the same as legal correctness. A model may predict a future event, but a court must justify a legal outcome through law, facts, procedure, discretion, and reasons. Kleinberg et al. show the promise of prediction in bail contexts, but their findings do not mean that prediction should replace adjudication [4]. Dressel and Farid show why courts should be sceptical of proprietary risk tools whose performance claims cannot be independently verified [5]. Green and Chen add that even accurate tools may interact with human judgment in unequal ways [6].

Bias is also more than a technical problem. Legal datasets may reflect policing patterns, litigation inequality, plea bargaining, settlement pressure, language barriers, and unequal access to representation. AI can encode these patterns and present them as objective. Krištofik's work is important because it shows that judicial AI can reproduce older institutional biases through new technology [16]. Responsible smart courts therefore need both technical bias audits and legal equality review.

Transparency must be practical. Full source-code disclosure may not always help litigants or judges. However, parties should know whether AI was used, for what purpose, what data categories were processed, what output was produced, what limitations were known, and how the human decision-maker treated the output. This form of procedural transparency is consistent with the NIST lifecycle approach and the EU AI Act's risk-based logic [10,14].

GenAI increases the need for clear boundaries. Unlike narrow decision-support tools, GenAI produces fluent legal-looking text. That fluency

creates a risk of false authority. Socol de la Osa and Remolina show that GenAI can improve access and efficiency but can also mislead judicial reasoning through hallucinated authorities and opaque reasoning pathways [7]. Legal-agent systems such as AgentsCourt demonstrate growing technical capacity to simulate court debate and generate judicial-style outputs [17]. Such tools may support training or research, but real adjudication requires stricter rules, verification, and disclosure.

The Indian experience supports the distinction between assistive and adjudicatory AI. Official materials describe AI-supported transcription, translation, e-filing scrutiny, legal research support, SUPACE, SUVAS, and NJDG as tools for efficiency and access, not replacements for judicial reasoning [11-13]. These uses are defensible when they remain transparent, auditable, and human-validated. They become more problematic if analytics are used to rank judges, pressure decision-makers, or predict outcomes without safeguards.

Public trust is not a public-relations matter. It is part of judicial legitimacy. Courts derive authority from independence, reasoned judgment, and the perception that parties were heard by a responsible decision-maker. Lopes notes that visible reliance on AI may affect perceptions of procedural fairness [19]. Sourdin and Barry similarly caution that AI governance must account for the wider roles of courts, not only case throughput [20]. A faster court that is less trusted has not delivered better justice.

Policy Implications

For judges, the central implication is that AI should remain an aid to understanding, organisation, and verification. Judges should not rely on unverified AI-generated summaries, citations, or recommendations when deciding rights and liabilities.

For court administrators, AI adoption should begin with process mapping. Courts should identify which tasks are administrative, informational, analytical, or adjudicatory before selecting tools. Low-risk uses such as translation

support or registry defect detection require different controls from high-risk uses such as risk scoring or judgment drafting.

For regulators and oversight bodies, the priority is to create public rules for judicial AI. These should include disclosure requirements, audit standards, equality-impact assessments, data-protection duties, incident reporting, and complaint pathways.

For vendors and procurement authorities, contracts should include audit rights, data-use limits, security obligations, model documentation, performance testing, human oversight requirements, and exit plans. Courts should avoid systems that cannot be independently evaluated or that create unclear responsibility.

Conclusion

AI can make smart courts more efficient and accessible, especially in filing support, registry scrutiny, transcription, translation, legal research, case management, document organisation, and data analytics. These uses can reduce delay, improve language access, and help courts handle large volumes of information. However, efficiency is not the same as justice. A court that is fast but opaque, biased, privacy-invasive, or unaccountable does not deliver digital justice.

The article concludes that AI should be used as a supportive tool, not as a replacement for judges. Judicial decision-making requires legal interpretation, factual evaluation, discretion, empathy, contextual judgment, public reasons, and accountability. AI can help organise information, but it cannot carry legal responsibility for the final decision. The correct model is AI-assisted justice under human judicial authority.

For responsible deployment, courts should adopt clear AI-use policies, independent validation, bias and equality-impact assessments, data-protection safeguards, audit logs, disclosure duties, contestability procedures, judicial training, and procurement standards. The future of digital justice should not be framed as a choice between human judges and machines. It should be framed as a constitutional design problem: how to use

computational tools while preserving the human, public, and reason-giving character of adjudication.

Limitations

This study is limited by its qualitative and review-based design. It does not include interviews with judges, lawyers, litigants, technologists, or court administrators. It does not test the performance of any AI tool, evaluate court datasets, or conduct a statistical meta-analysis. Because GenAI, court technology, and AI regulation are changing quickly, later developments may require updating the analysis.

The study also gives particular attention to India and selected international frameworks. Other jurisdictions may raise different issues because of different constitutional structures, data-protection laws, court cultures, and administrative models. A further limitation is that official institutional documents often describe intended use rather than actual practice. Future empirical work is therefore needed before the framework is generalised across jurisdictions.

Future Directions / Future Scope

Future research should move from normative analysis to empirical evaluation. Priority areas include courtroom studies of how judges and staff interact with AI tools, audits of AI-assisted translation and transcription accuracy, assessment of litigant understanding and procedural fairness, and validation of legal-research and summarisation tools against jurisdiction-specific legal corpora.

Courts should also develop public AI registers that identify each deployed tool, its purpose, data sources, oversight mechanism, vendor status, audit history, and complaint pathway. Future policy should develop standards for procurement, benchmark testing, model documentation, red-team review, privacy-preserving deployment, and human oversight. GenAI in legal research and judgment drafting requires special attention because legal fluency can create a false impression of reliability.

Declarations

Acknowledgment: The author acknowledges the published academic literature, institutional AI governance frameworks, and official court-technology materials that informed this review-based study.

Funding Statement: No external funding was received for this study.

Conflict of Interest Statement: The author declares no conflict of interest.

Ethics Statement: Ethical approval was not required because this article is based exclusively on publicly available literature, policy documents, regulatory materials, and official institutional sources. No human participants, personal data, or confidential court records were used.

References

- European Commission for the Efficiency of Justice. European ethical charter on the use of artificial intelligence in judicial systems and their environment. Strasbourg: Council of Europe; 2018.
- Sourdin T. Judge v robot? Artificial intelligence and judicial decision-making. *Univ N S W Law J.* 2018;41(4):1114-1133. doi:10.53637/ZGUX2213.
- Re RM, Solow-Niederman A. Developing artificially intelligent justice. *Stanford Technol Law Rev.* 2019;22:242-289.
- Kleinberg J, Lakkaraju H, Leskovec J, Ludwig J, Mullainathan S. Human decisions and machine predictions. *Q J Econ.* 2018;133(1):237-293. doi:10.1093/qje/qjx032.
- Dressel J, Farid H. The accuracy, fairness, and limits of predicting recidivism. *Sci Adv.* 2018;4(1):eaao5580. doi:10.1126/sciadv.aao5580.
- Green B, Chen Y. Disparate interactions: an algorithm-in-the-loop analysis of fairness in risk assessments. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*; 2019; Atlanta, GA. New York: Association for Computing Machinery; 2019. p. 90-99. doi:10.1145/3287560.3287563.
- Socol de la Osa DU, Remolina N. Artificial intelligence at the bench: legal and ethical challenges of informing-or misinforming-judicial decision-making through generative AI. *Data Policy.* 2024;6:e59. doi:10.1017/dap.2024.53.
- Organisation for Economic Co-operation and Development. OECD AI principles. Paris: OECD; 2024.
- United Nations Educational, Scientific and Cultural Organization. Recommendation on the ethics of artificial intelligence. Paris: UNESCO; 2021.
- European Parliament, Council of the European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. *Off J Eur Union.* 2024;L 2024/1689.
- Government of India, Ministry of Law and Justice. Use of AI in Supreme Court case management. New Delhi: Press Information Bureau; 2025 Mar 20.
- Government of India, Press Information Bureau. How AI is enhancing access to justice in India. New Delhi: Press Information Bureau; 2026 Feb 11.
- Supreme Court of India e-Committee. National Judicial Data Grid. New Delhi: eCourts Mission Mode Project; 2025.
- Tabassi E. Artificial Intelligence Risk Management Framework (AI RMF 1.0). Gaithersburg (MD): National Institute of Standards and Technology; 2023. Report No.: NIST AI 100-1. doi:10.6028/NIST.AI.100-1.
- Floridi L, Cowls J. A unified framework of five principles for AI in society. *Harv Data Sci Rev.* 2019;1(1). doi:10.1162/99608f92.8cd550d1.
- Krištofik A. Bias in AI (supported) decision making: old problems, new technologies. *Int J Court Admin.* 2025;16(1):3. doi:10.36745/ijca.598.
- He Z, Cao P, Wang C, Jin Z, Chen Y, Xu J, Li H, Liu K, Zhao J. AgentsCourt: building judicial decision-making agents with court debate simulation and legal knowledge augmentation. In: *Findings of the Association for*

Computational Linguistics: EMNLP 2024; 2024; Miami, FL. Association for Computational Linguistics; 2024. p. 9399-9416. doi:10.18653/v1/2024.findings-emnlp.549.

Završnik A. Algorithmic justice: algorithms and big data in criminal justice settings. *Eur J Criminol.* 2021;18(5):623-642. doi:10.1177/1477370819876762.

Lopes GP. Bias in adjudication and the promise of AI: challenges to procedural fairness. *Law Technol Hum.* 2025;7(1):47-67. doi:10.5204/lthj.3812.

Sourdin T, Barry B. AI and courts: challenges and opportunities. In: Zalnieriute M, Limante A, editors. *The Cambridge Handbook of AI and Technologies in Courts.* Cambridge: Cambridge University Press; 2026. p. 17-31. doi:10.1017/9781009744225.004.